arXiv:1512.06161v1 [cs.IT] 18 Dec 2015

On Locally Recoverable (LRC) Codes

Mario Blaum
IBM Almaden Research Center
San Jose, CA 95120

Abstract

We present simple constructions of optimal erasure-correcting LRC codes by exhibiting their parity-check matrices. When the number of local parities in a parity group plus the number of global parities is smaller than the size of the parity group, the constructed codes are optimal with a field of size at least the length of the code. We can reduce the size of the field to at least the size of the parity groups when the number of global parities equals the number of local parities in a parity group plus one.

Keywords: Erasure-correcting codes, Locally Recoverable codes, Reed-Solomon codes, Generalized Concatenated codes, Integrated Interleaved codes, Maximally Recoverable codes, MDS codes, Partial MDS (PMDS) codes, Sector-Disk (SD) codes, local and global parities, heavy parities.

1 Introduction

Erasure-correcting codes combining local and global properties have arisen considerable interest in recent literature (see for instance [2][7][10][11][14][16][17][18][19] and the references therein, a complete list of references on the subject is beyond the scope of this paper). The practical applications of these codes explain the reasons for this interest. For example, in storage applications involving multiple storage devices (like in the cloud), most of the failures involve only one device. In that case, it is convenient to recover the failed device "locally," that is, invoking a relatively small set of devices. This local recovery process mitigates the decrease in performance until the failed device is replaced. On the other hand, it is desirable to have extra protection (involving global parities that extend over all the devices) in case more complex failure events occur. In this case, performance will be impacted more severely (requiring even a halt in operations), but it is expected that this is a rare event and when it occurs, data loss is avoided. Examples of these types of applications are Microsoft Windows Azure Storage [11] and Xorbas built on top of Facebook's Hadoop system running HDFS-RAID [17].

A different type of application consists of extending RAID architectures in an array of storage devices to cover individual sector or page failures in each storage device in the array.

D	D	D	L	L
D	D	D	L	L
D	D	D	L	L
D	G	G	L	L
G	G	G	L	L
G	G	G	L	L

Figure 1: A 6×5 array showing the data symbols, $\ell = 2$ local parities per row and $g = 8$ global parities.

In effect, certain storage devices like flash memories decay in time and with use (number of writes). Every sector or page is protected by its own Error-Correcting Code (ECC), but an accumulation of errors may exceed the ECC capability of the sector. In that case, there is a "silent failure" that will not be discovered until the failed sector or page are accessed (scrubbing techniques are inconvenient for flash memory, since they may impact future performance). For example, in case an array of devices is protected using RAID 5, in the event a device fails, the sectors or pages in the failed device are recovered by XORing the corresponding sectors or pages in the remaining devices. However, if one of the sectors or pages in one of those remaining devices has failed (without a total device failure), data loss will occur. Adding an extra parity device (RAID 6) is expensive and moreover, may not correct some combinations of failures (like two sector or page failures in a row in addition to a device failure). This problem was studied in [2][4][15].

Next we give some notation and define mathematically the problem. Consider an erasure-correcting code $\mathcal{C}$ consisting of $m \times n$ arrays over a field $GF(q)$. Each row in an array is protected by ℓ local parities as part of an $[n, n-\ell, \ell+1]$ MDS code, i.e., up to ℓ erasures per row can be corrected. In addition, g global parities are added. Then we say that $\mathcal{C}$ is an $(m, n; \ell, g)$ Locally Recoverable (LRC) code. Code $\mathcal{C}$ has length mn and dimension $k = m(n-\ell) - g$. As vectors, we consider the coordinates of the arrays in $\mathcal{C}$ row-wise.

The situation is illustrated in Figure 1 for a 6×5 array with $\ell = 2$ local parities in each row and $g = 8$ global parities. So, this is a $[30, 10]$ code over a field $GF(q)$.

When an $(m, n; \ell, g)$ LRC code is desired, the key question is how to construct the g global parities in order to optimize the code. There are several criteria for optimization and they depend on the application considered. A weak optimization involves maximizing the minimum distance of the code. To that end, bounds on the minimum distance are needed.

Given two integers s and $t \geq 1$, denote by $\langle s \rangle_t$ the (unique) integer u, $0 \leq u \leq t-1$, such that $u \equiv s \pmod{t}$. For example, $\langle 5 \rangle_3 = 2$. Also, given x, denote by $\lfloor x \rfloor$ the floor of x. A Singleton type of bound on the minimum distance d of an $(m, n; \ell, g)$ LRC code is given by [7]

$$d \leq \ell + n \left\lfloor \frac{g}{n-\ell} \right\rfloor + \langle g \rangle_{n-\ell} + 1. \tag{1}$$

D	D	D	L	L
D	D	D	L	L
D	D	D	L	L
D	D	D	L	L
D	D	D	L	L
D	G	G	L	L

Figure 2: A 6×5 array with $\ell = 2$ and $g = 2$.

For example, in Figure 1, $\lfloor \frac{g}{n-\ell} \rfloor = 2$ and $\langle g \rangle_{n-\ell} = 2$, so, $d \leq 15$. This is easy to see since, if any three rows are erased, 15 symbols are lost, but we only have 14 parities to recover them, the 8 global parities and 6 local ones. The general argument to prove bound (1) proceeds similarly. We say that an LRC code is optimal when bound (1) is met with equality.

Let us point out that there are optimization criteria stronger than the minimum distance. The strongest one involves the so called Partial MDS (PMDS) codes [2][6] (in [6], PMDS codes are called Maximally Recoverable codes). We say that an $(m, n; \ell, g)$ LRC code is PMDS if, whenever each row of the code is punctured [12] in exactly ℓ locations, the punctured code is MDS with respect to the g parities (i.e., it can correct any g erasures after puncturing). A weaker requirement (but still stronger than simply meeting bound (1)) is provided by Sector-Disk (SD) codes [15]: these codes can correct any g erasures after ℓ columns in an array are erased (which corresponds to puncturing each row in ℓ locations, but these locations are fixed). Not surprisingly, PMDS and SD codes are harder to construct than optimal LRC codes. The known constructions require a large field [2][6], although in some cases efficient implementations (involving only XOR operations and rotations of vectors) were found when the field is defined by the irreducible polynomial $1 + x + x^2 + \cdots + x^{p-1}$, p a prime and $mn < p$ [2]. Also, some good constructions of PMDS and SD codes (involving fields of size at least $2mn$ and mn respectively) were found for $g = 2$ [4], but the general problem remains open.

When constructing optimal LRC codes, it is important to minimize the size of the field. Bound (1), being a Singleton type of bound, does not take into account this size. Most constructions of optimal LRC codes in literature involved fields of large size, but in [18], optimal LRC codes were obtained for fields of size $q \geq mn$, that is, the length of the code, like in the case of Reed-Solomon (RS) codes [12]. In fact, it was shown in [18] that regular RS codes are a special case of the construction (when $m = 1$). For a bound on LRC codes taking into account the size of the field, see [5].

In this paper, we examine a special case of $(m, n; \ell, g)$ LRC codes that is important in applications [11][17]: the case $\ell + g < n$, hence, $\lfloor \frac{g}{n-\ell} \rfloor = 0$ and $\langle g \rangle_{n-\ell} = g$, so bound (1) becomes

$$d \leq \ell + g + 1. \tag{2}$$

For example, in Figure 2, we have such a situation for a 6×5 array with $\ell = 2$ and $g = 2$,

thus bound (2) gives $d \leq 5$. The assumption $\ell + g < n$ allows us to obtain some very simple constructions.

The paper is organized as follows: in Section 2, in addition to the condition $\ell + g < n$, we consider $(m, n; \ell, g)$ LRC codes under the assumption $g \leq \ell + 1$. Then we use the techniques of Generalized Concatenated codes [3] to obtain optimal $(m, n; \ell, g)$ LRC codes over fields of size $q \geq n$ (let us mention that there exist similar constructions in literature for other applications, like magnetic recording [1][8][9][13]). In Section 3 we present a simple construction over a field of size $q \geq mn$ (as in [18]) that had been given in [2] as a candidate for PMDS codes, and we show that the construction gives optimal $(m, n; \ell, g)$ LRC codes when $\ell + g < n$. Finally, in Section 4 we draw some conclusions justifying the assumption $\ell + g < n$.

2 Construction of Optimal LRC Codes Using Generalized Concatenation Techniques

Denote by $RS(n, r; i, j)$, where $0 \leq i$, $0 \leq j$ and $1 \leq r < n$, the following parity-check matrix, corresponding to a Reed-Solomon (RS) code of length n and r parities:

$$RS(n,r;i,j) = \begin{pmatrix} \alpha^{ij} & \alpha^{i(j+1)} & \alpha^{i(j+2)} & \dots & \alpha^{i(j+n-1)} \\ \alpha^{(i+1)j} & \alpha^{(i+1)(j+1)} & \alpha^{(i+1)(j+2)} & \dots & \alpha^{(i+1)(j+n-1)} \\ \alpha^{(i+2)j} & \alpha^{(i+2)(j+1)} & \alpha^{(i+2)(j+2)} & \dots & \alpha^{(i+2)(j+n-1)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \alpha^{(i+r-1)j} & \alpha^{(i+r-1)(j+1)} & \alpha^{(i+r-1)(j+2)} & \dots & \alpha^{(i+r-1)(j+n-1)} \end{pmatrix}. \quad (3)$$

Also, denote by I_m the $m \times m$ identity matrix and by $A \otimes B$ the Kronecker product between matrices A and B [12].

Consider the $(m, n; \ell, g)$ LRC code over a field $GF(q)$ given by the $(\ell m + g) \times mn$ matrix [3]

$$H = \begin{pmatrix} I_m \otimes RS(n, \ell; 0, 0) \\ \hline \overbrace{(1\ 1\ \dots\ 1)}^{m} \otimes RS(n, g; \ell, 0) \end{pmatrix}. \quad (4)$$

Then we have the following lemma:

Lemma 2.1 Let $g \leq \ell + 1$, $\ell + g < n$ and consider a field $GF(q)$ such that $q > n$. Then the $(m, n; \ell, g)$ LRC code over $GF(q)$ whose parity-check matrix H is given by (4) is an optimal LRC code.

Proof: Assume that $g = \ell + 1$ (the cases $g \leq \ell$ are similar). We show that the code can correct any $g + \ell = 2\ell + 1$ erasures. In effect, assuming that $2\ell + 1$ erasures have occurred,

first we correct all the rows having up to ℓ erasures using the ℓ local parities in those rows. If any erasures are left after this process, they are all in one row with at most $2\ell+1$ erasures. But from (4), there are $2\ell+1$ parities corresponding to a RS code (ℓ local parities and $\ell+1$ global parities) to correct these at most $2\ell+1$ erasures, so all the erasures can be corrected since they involve inverting a $(2\ell+1)\times(2\ell+1)$ Vandermonde matrix over the field $GF(q)$: since $2\ell+1=\ell+g<n<q$, the elements in the Vandermonde matrix are distinct and the matrix is invertible. Thus, the minimum distance satisfies $d\geq\ell+g+1$. But, since $\ell+g<n$, bound (2) gives $d=\ell+g+1$ and the code is optimal.

□

Example 2.1 Take $m=3$, $n=6$, $\ell=2$ and $g=3$, then the parity-check matrix over $GF(8)$ given by (4) is (notice, $\alpha^7=1$)

$$H = \left(\begin{array}{cccccc|cccccc|cccccc}
1 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & \alpha & \alpha^2 & \alpha^3 & \alpha^4 & \alpha^5 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & \alpha & \alpha^2 & \alpha^3 & \alpha^4 & \alpha^5 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 1 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & \alpha & \alpha^2 & \alpha^3 & \alpha^4 & \alpha^5 \\
\hline
1 & \alpha^2 & \alpha^4 & \alpha^6 & \alpha & \alpha^3 & 1 & \alpha^2 & \alpha^4 & \alpha^6 & \alpha & \alpha^3 & 1 & \alpha^2 & \alpha^4 & \alpha^6 & \alpha & \alpha^3 \\
1 & \alpha^3 & \alpha^6 & \alpha^2 & \alpha^5 & \alpha & 1 & \alpha^3 & \alpha^6 & \alpha^2 & \alpha^5 & \alpha & 1 & \alpha^3 & \alpha^6 & \alpha^2 & \alpha^5 & \alpha \\
1 & \alpha^4 & \alpha & \alpha^5 & \alpha^2 & \alpha^6 & 1 & \alpha^4 & \alpha & \alpha^5 & \alpha^2 & \alpha^6 & 1 & \alpha^4 & \alpha & \alpha^5 & \alpha^2 & \alpha^6
\end{array}\right).$$

It is easy to see directly that the parity-check matrix above allows for the correction of any 5 erasures, so the minimum distance is 6 and the code is optimal.

□

We can obtain an optimal $(m,n;\ell,g)$ LRC code also for $q=n$ by using extended RS codes, which requires a slight modification of parity-check matrix (4). In effect, let

$$ERS(n,r;j) = \left(\begin{array}{ccccccc}
1 & 1 & 1 & \dots & 1 & 1 \\
\alpha^{j} & \alpha^{j+1} & \alpha^{j+2} & \dots & \alpha^{j+n-2} & 0 \\
\alpha^{2j} & \alpha^{2(j+1)} & \alpha^{2(j+2)} & \dots & \alpha^{2(j+n-2)} & 0 \\
\alpha^{3j} & \alpha^{3(j+1)} & \alpha^{3(j+2)} & \dots & \alpha^{3(j+n-2)} & 0 \\
\vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\
\alpha^{(r-1)j} & \alpha^{(r-1)(j+1)} & \alpha^{(r-1)(j+2)} & \dots & \alpha^{(r-1)(j+n-2)} & 0
\end{array}\right) \qquad (5)$$

and consider the $(m,n;\ell,g)$ LRC code over a field $GF(q)$ given by the $(\ell m+g)\times mn$ matrix

$$EH \;=\; \left(\frac{I_m \otimes ERS(n,\ell;0)}{(\overbrace{1\ 1\ \ldots\ 1}^{m}) \;\otimes\; \left(RS(n-1,g;\ell,0) \;\mid\; \underline{0}_{g\times 1} \right)} \right). \tag{6}$$

where $\underline{0}_{i\times j}$ denotes an $i \times j$ zero matrix.

Lemma 2.2 Let $g \leq \ell + 1$, $\ell + g < n$ and take a field $GF(q)$ such that $q = n$. Then the $(m, n; \ell, g)$ LRC code over $GF(q)$ whose parity-check matrix EH is given by (6) is an optimal $(m, n; \ell, g)$ LRC code.

Proof: Completely analogous to the proof of Lemma 2.1, since an extended RS code is also MDS.

□

Example 2.2 Consider the situation described in [17]. There, the authors present a $(2, 8; 1, 4)$ LRC code over $GF(16)$ with minimum distance 5. Since $\ell + g = 5$, this code is not optimal. If we construct the parity-check matrix EH as given by (6) for $m = 2$, $n = 8$, $\ell = 2$ and $g = 2$ over $GF(8)$, we obtain

$$EH \;=\; \left(\begin{array}{cccccccc|cccccccc}
1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & \alpha & \alpha^2 & \alpha^3 & \alpha^4 & \alpha^5 & \alpha^6 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & \alpha & \alpha^2 & \alpha^3 & \alpha^4 & \alpha^5 & \alpha^6 & 0 \\
\hline
1 & \alpha^2 & \alpha^4 & \alpha^6 & \alpha & \alpha^3 & \alpha^5 & 0 & 1 & \alpha^2 & \alpha^4 & \alpha^6 & \alpha & \alpha^3 & \alpha^5 & 0 \\
1 & \alpha^3 & \alpha^6 & \alpha^2 & \alpha^5 & \alpha & \alpha^4 & 0 & 1 & \alpha^3 & \alpha^6 & \alpha^2 & \alpha^5 & \alpha & \alpha^4 & 0
\end{array}\right).$$

By Lemma 2.2, this code has minimum distance 5 and is an optimal LRC code. Compared to the code in [17], the minimum distance is the same, but this code can correct up to two erasures per row and it requires a smaller field, i.e., $GF(8)$ as opposed to $GF(16)$. However, the locality in [17], that is, the number of symbols necessary to reconstruct a lost (data) symbol, is 5, while in this example it is 6. Moreover, it was proven in [17] that if the locality is 5, the minimum distance is at most 5, so the code in [17] satisfies a different type of optimality. The decision on which code to use and which type of optimality is preferable depends on the application, there are trade-offs that need to be considered.

□

3 Construction of Optimal LRC Codes when $\ell + g < n$

Consider Construction 3.2 in [2]. For the sake of completeness, we present it next. Let the $(m, n; \ell, g)$ LRC code over $GF(q)$, $q > mn$, be given by the parity-check matrix

$$H(m,n;\ell,g) = \left(\begin{array}{c|c|c|c} RS(n,\ell;0,0) & \underline{0}_{\ell\times n} & \cdots & \underline{0}_{\ell\times n} \\ \underline{0}_{\ell\times n} & RS(n,\ell;0,r) & \cdots & \underline{0}_{\ell\times n} \\ \vdots & \vdots & \ddots & \vdots \\ \underline{0}_{\ell\times n} & \underline{0}_{\ell\times n} & \cdots & RS(n,\ell;0,(m-1)\ell) \\ \hline \multicolumn{4}{c}{RS(mn,g;\ell,0)} \end{array}\right), \quad (7)$$

Example 3.1 Take $m=3$, $n=5$, $\ell=1$, $g=3$ and the field $GF(16)$. According to (7), we obtain

$$H(3,5;1,3) = \left(\begin{array}{ccccc|ccccc|ccccc} 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 \\ \hline 1 & \alpha & \alpha^2 & \alpha^3 & \alpha^4 & \alpha^5 & \alpha^6 & \alpha^7 & \alpha^8 & \alpha^9 & \alpha^{10} & \alpha^{11} & \alpha^{12} & \alpha^{13} & \alpha^{14} \\ 1 & \alpha^2 & \alpha^4 & \alpha^6 & \alpha^8 & \alpha^{10} & \alpha^{12} & \alpha^{14} & \alpha & \alpha^3 & \alpha^5 & \alpha^7 & \alpha^9 & \alpha^{11} & \alpha^{13} \\ 1 & \alpha^3 & \alpha^6 & \alpha^9 & \alpha^{12} & 1 & \alpha^3 & \alpha^6 & \alpha^9 & \alpha^{12} & 1 & \alpha^3 & \alpha^6 & \alpha^9 & \alpha^{12} \end{array}\right).$$

Similarly, if $m=3$, $n=5$, $\ell=2$ and $g=2$, (7) gives

$$H(3,5;2,2) = \left(\begin{array}{ccccc|ccccc|ccccc} 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & \alpha & \alpha^2 & \alpha^3 & \alpha^4 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \alpha^5 & \alpha^6 & \alpha^7 & \alpha^8 & \alpha^9 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \alpha^{10} & \alpha^{11} & \alpha^{12} & \alpha^{13} & \alpha^{14} \\ \hline 1 & \alpha^2 & \alpha^4 & \alpha^6 & \alpha^8 & \alpha^{10} & \alpha^{12} & \alpha^{14} & \alpha & \alpha^3 & \alpha^5 & \alpha^7 & \alpha^9 & \alpha^{11} & \alpha^{13} \\ 1 & \alpha^3 & \alpha^6 & \alpha^9 & \alpha^{12} & 1 & \alpha^3 & \alpha^6 & \alpha^9 & \alpha^{12} & 1 & \alpha^3 & \alpha^6 & \alpha^9 & \alpha^{12} \end{array}\right).$$

□

We have the following lemma:

Lemma 3.1 Consider the $(m, n; \ell, g)$ LRC code over $GF(q)$ with $q > mn$ whose parity-check matrix $H(m, n; \ell, g)$ is given by (7). Then the code has minimum distance $d = \ell + g + 1$.

Proof: Notice that if we take the linear combination of the rows of $H(m,n;\ell,g)$ consisting of XORing rows $0,\ell,2\ell,\ldots,(m-1)\ell$, then rows $1,\ell+1,2\ell+1,\ldots,(m-1)\ell+1$, and so on, up to rows $\ell-1,2\ell-1,3\ell-1,\ldots,m\ell-1$, we obtain the $(\ell+g)\times mn$ matrix $RS(mn,\ell+g;0,0)$.

Consider any $(\ell m+g)\times(\ell+g)$ submatrix of $H(m,n;\ell,g)$. We have to prove that this matrix has rank $\ell+g$, which is at least the rank of any linear combination of its rows. In particular, it is at least the rank of the corresponding $(\ell+g)\times(\ell+g)$ submatrix of $RS(mn,\ell+g;0,0)$, which is $\ell+g$ since it is a Vandermonde matrix and $mn<q$.

□

Corollary 3.1 Consider an $(m,n;\ell,g)$ LRC code with the conditions of Lemma 3.1 and assume that $\ell+g<n$. Then the code is an optimal LRC code.

Proof: Simply notice that, by Lemma 3.1, inequality (2) is met with equality.

□

Example 3.2 Consider the $(3,5;2,2)$ LRC code whose parity-check matrix is $H(3,5;2,2)$ as given in Example 3.1. According to Lemma 3.1, its minimum distance is 5. To view this, XORing rows 0, 2 and 4, and then rows 1, 3 and 5 of $H(3,5;2,2)$, the 4×15 matrix

$$RS(15,4;0,0) = \left(\begin{array}{ccccc|ccccc|ccccc} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & \alpha & \alpha^2 & \alpha^3 & \alpha^4 & \alpha^5 & \alpha^6 & \alpha^7 & \alpha^8 & \alpha^9 & \alpha^{10} & \alpha^{11} & \alpha^{12} & \alpha^{13} & \alpha^{14} \\ \hline 1 & \alpha^2 & \alpha^4 & \alpha^6 & \alpha^8 & \alpha^{10} & \alpha^{12} & \alpha^{14} & \alpha & \alpha^3 & \alpha^5 & \alpha^7 & \alpha^9 & \alpha^{11} & \alpha^{13} \\ 1 & \alpha^3 & \alpha^6 & \alpha^9 & \alpha^{12} & 1 & \alpha^3 & \alpha^6 & \alpha^9 & \alpha^{12} & 1 & \alpha^3 & \alpha^6 & \alpha^9 & \alpha^{12} \end{array}\right)$$

is obtained. Any 8×4 submatrix of $H(3,5;2,2)$ has rank 4, since the corresponding 4×4 submatrix of $RS(15,4;0,0)$ also has rank 4.

□

In Lemma 3.1, we have shown how to construct $(m,n;\ell,g)$ LRC codes over $GF(q)$ with $q>mn$ and minimum distance $d=\ell+g+1$. We can also construct a code with these parameters and $q=mn$ by considering extended RS codes as done in Section 2. Using (5) consider the $(m,n;\ell,g)$ LRC code over $GF(q)$, $q=mn$, given by the parity-check matrix

$$EH(m,n;\ell,g) = \left(\begin{array}{c|c|c|c} RS(n,\ell;0,0) & \underline{0}_{\ell\times n} & \cdots & \underline{0}_{\ell\times n} \\ \underline{0}_{\ell\times n} & RS(n,\ell;0,r) & \cdots & \underline{0}_{\ell\times n} \\ \vdots & \vdots & \ddots & \vdots \\ \underline{0}_{\ell\times n} & \underline{0}_{\ell\times n} & \cdots & ERS(n,\ell;(m-1)\ell) \\ \hline \multicolumn{3}{c|}{RS(mn-1,g;\ell,0)} & \underline{0}_{g\times 1} \end{array}\right). \quad (8)$$

Lemma 3.2 The $(m, n; \ell, g)$ LRC code over $GF(q)$ with $q = mn$ whose parity-check matrix $EH(m, n; \ell, g)$ is given by (8) has minimum distance $d = \ell + g + 1$.

Proof: Completely analogous to the proof of Lemma 3.1, except that after making the XORs described there we obtain the $(\ell + g) \times mn$ matrix $ERS(mn, \ell + g; 0)$, corresponding to an extended RS code and which is MDS.

□

Example 3.3 Consider again the situation of a $(2, 8; 1, 4)$ LRC code over $GF(16)$ with minimum distance 5 described in [17]. If we construct $EH(2, 8; 1, 4)$ also over $GF(16)$ as given by (8), we obtain

$$EH(2,8;1,4) = \left(\begin{array}{cccccccc|cccccccc} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ \hline 1 & \alpha & \alpha^2 & \alpha^3 & \alpha^4 & \alpha^5 & \alpha^6 & \alpha^7 & \alpha^8 & \alpha^9 & \alpha^{10} & \alpha^{11} & \alpha^{12} & \alpha^{13} & \alpha^{14} & 0 \\ 1 & \alpha^2 & \alpha^4 & \alpha^6 & \alpha^8 & \alpha^{10} & \alpha^{12} & \alpha^{14} & \alpha & \alpha^3 & \alpha^5 & \alpha^7 & \alpha^9 & \alpha^{11} & \alpha^{13} & 0 \\ 1 & \alpha^3 & \alpha^6 & \alpha^9 & \alpha^{12} & 1 & \alpha^3 & \alpha^6 & \alpha^9 & \alpha^{12} & 1 & \alpha^3 & \alpha^6 & \alpha^9 & \alpha^{12} & 0 \\ 1 & \alpha^4 & \alpha^8 & \alpha^{12} & \alpha & \alpha^5 & \alpha^9 & \alpha^{13} & \alpha^2 & \alpha^6 & \alpha^{10} & \alpha^{15} & \alpha^3 & \alpha^7 & \alpha^{11} & 0 \end{array}\right).$$

By Lemma 3.2, this code has minimum distance 6 and hence is an optimal LRC code. The locality, though, is 7, as opposed to 5 in [17].

□

4 Conclusions

We have studied the problem of $(m, n; \ell, g)$ LRC codes and presented some simple constructions for the case $\ell + g < n$. The size of the field is $q \geq mn$ and in some cases (when in addition, $g \leq \ell + 1$), $q \geq n$.

The general case without the restriction $\ell + g < n$ is handled in [18] and the construction there is much more involved. For example, consider 6×5 arrays like in Figures 1 and 2. In Figure 2 with $\ell = g = 2$, an optimal LRC code with minimum distance $d = 5$ is obtained with the construction presented. However, if $g = 3$, the minimum distance is $d = 6$, while the optimal, according to bound (1), is $d = 7$. The remarkable construction in [18] effectively provides such a code. So, the question is, how relevant is the condition $\ell + g < n$ in applications?

A channel requiring a large number of global parities g may neglect the impact of local recovery. As stated in the Introduction, it is desirable that the local parities are the ones most heavily used. The reason they have been introduced is due to a performance issue: they allow for fast recovery in case the number of erasures in a row does not exceed ℓ. The

global parities are used in case the erasure-correcting power of the local parities is exceeded, but for the system to be efficient, these should be relatively rare cases. If the global parities need to be frequently invoked, we may do better with an MDS code (saving in parity), or, with the same amount of parity, having a much stronger code. For example, if we take the $[30, 10]$ code of Figure 1, in the optimal case we can correct 14 erasures. But if we used an MDS code we could correct 20 erasures by sacrificing the locality. Alternatively, we could use a $[30, 16]$ MDS code that can also correct up to 14 erasures, but with a much better rate.

References

[1] K. Abdel-Ghaffar and M. Hassner, "Multilevel codes for data storage channels," IEEE Trans. on Information Theory, vol. IT-37, pp. 735–41, May 1991.

[2] M. Blaum, J. L. Hafner and S. Hetzler, "Partial-MDS Codes and their Application to RAID Type of Architectures," IEEE Trans. on Information Theory, vol. IT-59, pp. 4510-19, July 2013.

[3] M. Blaum and S. Hetzler, "Generalized Concatenated Types of Codes for Erasure Correction," arXiv:1406.6270, June 2014.

[4] M. Blaum, J. S. Plank, M. Schwartz and E. Yaakobi, "Partial MDS (PMDS) and Sector-Disk (SD) Codes that Tolerate the Erasure of Two Random Sectors," IEEE International Symposium on Information Theory (ISIT'14), Honolulu, Hawaii, June 2014.

[5] V. Cadambe and A. Mazumdar, "An Upper Bound on the Size of Locally Recoverable Codes," Proc. IEEE Symposium on Network Coding, pp. 1–5, June 2013.

[6] P. Gopalan, C. Huang, B. Jenkins and S. Yekhanin, "Explicit Maximally Recoverable Codes with Locality," IEEE Trans. on Information Theory, vol. IT-60, pp. 5245–56, September 2014.

[7] P. Gopalan, C. Huang, H. Simitci and S. Yekhanin, "On the Locality of Codeword Symbols," IEEE Trans. on Information Theory, vol. IT-58, pp. 6925–34, November 2012.

[8] J. Han and L. A. Lastras-Montaño, "Reliable Memories with Subline Accesses," ISIT 2007, IEEE International Symposium on Information Theory, pp. 2531–35, June 2007.

[9] M. Hassner, K. Abdel-Ghaffar, A. Patel, R. Koetter and B. Trager, "Integrated Interleaving – A Novel ECC Architecture," IEEE Transactions on Magnetics, Vol. 37, No. 2, pp. 773–5, March 2001.

[10] C. Huang, M. Chen and J. Li, "Erasure-Resilient Codes Having Multiple Protection Groups," US Patent 7,930,611, April 2011.

[11] C. Huang, H. Simitci, Y. Xu, A. Ogus, B. Calder, P. Gopalan, J. Li and S. Yekhanin, "Erasure Coding in Windows Azure Storage," 2012 USENIX Annual Technical Conference, Boston, Massachussetts, June 2012.

[12] F. J. MacWilliams and N. J. A. Sloane, "The Theory of Error-Correcting Codes," North Holland, Amsterdam, 1977.

[13] A. Patel, "Two-Level Coding for Error-Control in Magnetic Disk Storage Products," IBM Journal of Research and Development, vol. 33, pp. 470–84, 1989.

[14] D. S. Papailiopoulos and A. G. Dimakis, "Locally Repairable Codes," IEEE Trans. on Information Theory, vol.IT-60, pp. 5843-55, October 2014.

[15] J. S. Plank and M. Blaum, "Sector-Disk (SD) Erasure Codes for Mixed Failure Modes in RAID Systems," ACM Transactions on Storage, Vol. 10, No. 1, Article 4, January 2014.

[16] A. S. Rawat, O. O. Koyluoglu, N. Silberstein and S. Vishwanath, "Optimal Locally Repairable and Secure Codes for Distributed Storage Systems," IEEE Trans. on Information Theory, vol.IT-60, pp. 212-36, January 2014.

[17] M. Sathiamoorthy, M. Asteris, D. Papailiopoulos, A. G. Dimakis, R. Vadali, S. Chen, and D. Borthakur, "XORing Elephants: Novel Erasure Codes for Big Data," Proceedings of VLDB, Vol. 6, No. 5, pp. 325–336, August 2013.

[18] I. Tamo and A. Barg, "A Family of Optimal Locally Recoverable Codes," IEEE Trans. on Information Theory, vol. IT-60, pp. 4661–76, August 2014.

[19] I. Tamo, Z. Wang and J. Bruck, "Zigzag Codes: MDS Array Codes With Optimal Rebuilding," IEEE Trans. on Information Theory, vol. IT-59, pp. 1597–616, March 2013.